

A UNIFIED MODEL OF RULE-SET LEARNING AND SELECTION

by


Pierson Fleischer

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy



Department of Psychological Sciences

West Lafayette, Indiana

December 2018

**THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL**

Dr. Sébastien Hélie, Chair

Department of Psychological Sciences

Dr. Gregory S. Francis

Department of Psychological Sciences

Dr. Richard Schweickert

Department of Psychological Sciences

Dr. Shawn W. Ell

Department of Psychology, University of Maine

Approved by:

Dr. David Rollock

Head of the Graduate Program

TABLE OF CONTENTS

LIST OF TABLES	5
LIST OF FIGURES	7
ABSTRACT	12
INTRODUCTION	13
Rules	13
Physiological Basis for Rules	15
Behavioral Investigations of Rule-Set Switching	16
Summary and Outline	22
A NEW THEORY OF RULE-SET REPRESENTATION AND LEARNING	24
Feedback and Learning in the PFC	25
Presynaptic Inhibitory Gating	28
Rule-Set Maintenance in Working Memory	29
Hierarchical Organization of the Prefrontal Cortex	33
The Theory	35
THE MODEL	38
Cortical Pyramidal Cells	38
Non-Cortical Cells	40
Modular Design	42
Response Selection Module	42
Cued Rule-Set Selection	46
Default Model Architecture and Input Parameters	48
Summary	53
EXPERIMENT 1 – BADRE, KAYSER, AND D’ESPOSITO 2010	65
Task	65
Human Data	67
Task-Specific Model Parameters	67
Results	72
Discussion	74
EXPERIMENT 2 – COLLINS AND FRANK 2013	75

Task 75

Human Data 77

Task-Specific Model Parameters. 78

Model Results 80

Discussion 80

EXPERIMENT 3 – ALPORT, STYLES, AND HSIEH 1994 86

Task 86

Human Data 87

Task-Specific Model Parameters. 87

Model Results 90

Discussion 90

EXPERIMENT 4 – VAN ‘T WOUT, LAVRIC, AND MONSELL 2015 96

Task 96

Human Data 97

Task-Specific Model Parameters. 97

Model Results 100

Discussion 100

CONCLUSIONS 106

Key Features 106

Accomplishments. 107

Implications and Predictions 108

Extensions, Improvements, and Future Work. 109

LIST OF REFERENCES 112

LIST OF TABLES

Table 1: Fixed Connection Weights in the Model for the Default Model Implementation	52
Table 2: Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell and the Stimulus Cell Outputs in the Model for the Badre, Kaysre and D’Esposito Tasks	70
Table 3: Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell and the Internal Cue Cell Outputs in the Model for the Badre, Kaysre and D’Esposito Tasks.	71
Table 4: Fixed Connection Weights in the Model for the Collins and Frank Tasks That Differ From the Default Values.	81
Table 5: Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell on the Stimulus Cell Outputs in the Model for the Collins and Frank Tasks	82
Table 6: Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell on the Internal Cue Cell Outputs in the Model for the Collins and Frank Tasks	83
Table 7: Fixed Connection Weights in the Model for the Alport, Styles, and Hsieh Tasks That Differ From the Default Values.	91
Table 8: Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell on the Stimulus Cell Outputs in the Model for the Alport, Styles, and Hsieh Tasks.	92
Table 9: Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell on the Internal Cue Cell Outputs in the Model for the Alport, Styles, and Hsieh Tasks.	93
Table 10: Fixed Connection Weights in the Model for the van ‘t Wout, Lavric, and Monsell Task That Differ From the Default Values.	101

Table 11: Parameter Values for the Plastic Non-Somatic Connections Between the Concrete Rule-Set Cell on the Stimulus Cell Outputs in the Model for the van 't Wout, Lavric, and Monsell Task	102
Table 12: Parameter Values for the Plastic Non-Somatic Connections Between the Abstract Rule-Set Cell on the Internal Cue Cell Outputs in the Model for the van 't Wout, Lavric, and Monsell Task	103

LIST OF FIGURES

<p>Figure 1: The FROST model architecture showing the pathway through the basal ganglia and the parallel maintenance loops between the PFC and the thalamus (from Ashby et al., 2005)</p>	31
<p>Figure 2: From Koechlin, Ody, & Kouneiher 2003. fMRI results showing the hierarchical organization of the PFC. Areas in green were more active during tasks with more possible responses (large concrete rule-sets). Areas in yellow were more active during tasks with more possible cues – each corresponding to a different rule-set (many concrete rule-sets). Areas in red were more active in tasks where the stimulus-response and cue-rule-set associations conflicted with a larger percentage of other task blocks (uncommon abstract rule-set)</p>	34
<p>Figure 3: The broad connectivity diagram of the proposed theory. Blue connections represent excitatory connections and red connections represent inhibitory connections. Purple connections are used only with the feedback area (orbitomedial PFC) and represent both excitatory and inhibitory connections. Numbered areas are described in the text</p>	37
<p>Figure 4: The connectivity diagram of the response selection model segment. Non-somatic connections are denoted using a gray field over the receiving somatic connection.</p>	43
<p>Figure 5: The connectivity diagram of the response selection model segment. Non-somatic connections are denoted using a gray field over the receiving somatic connection.</p>	47

Figure 6: A connectivity diagram including spike trains for the major brain areas included in the model. Blue connections are excitatory and red connections are inhibitory. Dotted lines are plastic while solid lines are fixed. Spike train data was taken from a trial from the task from van 't Wout, Lavric, & Monsell (2015) (see Chapter 7 for a full description). The spike train shown as from the most active cell in each brain area and the graphs are cut off at 2500ms since there is no change in activity after that point. 50

Figure 7: We assume a rule-set has already been selected (see other figure for rule-set selection process). The selected rule-set cell sends strong, blanket non-somatic inhibition to the outgoing connections of stimulus cells in the irrelevant dimension (Dim 2) and selective, milder non-somatic inhibition to the outgoing connections of stimulus cells in the relevant dimension (Dim 1) 54

Figure 8: The stimulus is presented and the stimulus cells in each dimension that correspond to the stimulus's features receive excitatory input and become active 55

Figure 9: The strong inhibition on the outgoing connections from stimulus cells of the irrelevant dimension prevents the active cell in that dimension from sending any signal to the PMC cells. Assume some learning has taken place in the relevant dimension and that the inhibition is stronger on the connection to the left PMC cell. Thus, the signal that the right PMC cell receives is stronger than what the left receives. 56

Figure 10: Both of the PMC cells become active and begin laterally inhibiting the incoming connections to each other. However, the weaker signal received by the left PMC cell results in less activity and less lateral inhibition compared to the cell on the right. If this were early in the training, the excitation received by the PMC cells would be nearly equal. In this case, the necessary differences in activity would arise primarily through noise 57

- Figure 11: The weak inhibition from the left PMC cell has only a small effect on the input and, by extension, the activity of the right PMC cell. Meanwhile, the strong inhibition from the right PMC cell stifles what little excitatory input the left PMC cell received. The right cell is now free of the left's inhibition and there is no possible way for the left or any other PMC cell to become active. The output of the right cell accumulates until it reaches the response threshold and the model responds accordingly. 58
- Figure 12: In rule-set selection there is a lot of holdover from the previous trial and therefore a lot more going on at the start. The abstract rule-set cell is active and gating the connections from the cue cells to the concrete rule-set cells (although the cue cells themselves are not active). The concrete rule-set cell that was used in the previous trial stays active into the next and its activity is maintained through a pair of excitatory connections with a thalamic cell. The external segment of the globus pallidus is naturally active even in the absence of excitatory input. The GPe inhibits the GPi which otherwise would be similarly active 59
- Figure 13: Once the trial begins, an excitatory signal is sent to the GPi from the subthalamic nucleus in order to clear the previous rule-set from working memory. This excitation counteracts the inhibition from the GPe and the GPi becomes active. The GPi begins to inhibit the thalamus 60
- Figure 14: The inhibition from the GPi prevents the thalamus from firing. The input to the concrete rule-set cell dwindles and soon both it and the thalamus will be completely inactive 61
- Figure 15: For brevity, assume that the cue is presented at about the same time that the reset signal stops. Without the excitation of the reset signal the GPi begins to go quiet. The cue cell representing the cue becomes active and begins sending its signal, gated by the abstract rule-set cell. Once again, assume that some learning has taken place and that the gating inhibition is uneven. 62

- Figure 16: Both of the concrete rule-set cells become active and begin exciting their corresponding thalamus cells as well as laterally inhibiting the incoming connections to each other. However, the weaker signal received by the right concrete rule-set cell results in less activity and less lateral inhibition compared to the cell on the left. If this were early in the training, the excitation received by the concrete rule-set cells would be nearly equal. In this case, the necessary differences in activity would arise primarily through noise. 63
- Figure 17: The weak inhibition from the right concrete rule-set cell has only a small effect on the input and, by extension, the activity of the left concrete rule-set cell, especially once it begins reverberating with its paired thalamus cell. Meanwhile, the strong inhibition from the left concrete rule-set cell stifles what little excitatory input the right concrete rule-set cell received. What little excitation makes it to the thalamus and back is not enough to start a positive feedback loop. The right cell is now free of the left's inhibition and there is no possible way for the right or any other concrete rule-set cell to become active. The response selection module of the model is now gated solely by the left concrete rule-set cell 64
- Figure 18: A sample of the stimuli from Badre, Kayser, and D'Esposito's (2010). Note how the hierarchical stimuli can be arranged by shape for red-bordered stimuli and by orientation for blue-bordered stimuli while no such pattern exists for the flat stimuli. Figures taken from Badre Kayser, and D'Esposito's (2010) 66
- Figure 19: Human results from Badre, Kayser, and D'Esposito's (2010). Graphs show the learning curve estimates and 90% confidence interval of the most typical subject for each condition. Graphs taken from Badre Kayser, and D'Esposito's (2010) 68
- Figure 20: Model data from performing the task from Bardre, Kayser, and D'Esposito (2010). 73

- Figure 21: A) Stimulus-response (responses denoted as A_1 , A_2 , etc.) pairings for the Collins and Frank task. Note how yellow and blue stimuli have the same response pattern with regards to shape while green stimuli have a unique response pattern. B) The task representation that uses shape to select a task set (TS) and color to select a response. Note how in phase 2 response associations for blue and green stimuli are added to the existing task-sets. No transfer can take place. C) The task representation that uses color to select a task set (TS) and shape to select a response. Note how in phase 2 new task-sets are used for blue and green stimuli allowing blue stimuli to use the task-set already created for yellow stimuli. This enables transfer. Figures taken from Collins and Frank (2013) 76
- Figure 22: Human results for the third of subjects with the highest color-switch cost minus shape-shift cost (Group 1) and the third of subjects with the lowest color-switch cost minus shape-shift cost (Group 3). The graphs show average accuracy (error bars are standard error) and the inserts show green errors minus blue errors by error type: errors made due to neglecting color (NC), errors made due to neglecting shape (NS), and errors made due to neglecting both color and shape (NA). Figures taken from Collins and Frank (2013). 79
- Figure 23: Model performance in the Collins and Frank task. The graphs show average accuracy and the inserts show green errors minus blue errors by error type: errors made due to neglecting color (NC), errors made due to neglecting shape (NS), and errors made due to neglecting both color and shape (NA) 84
- Figure 24: Human data from Alport, Styles, and Hsieh (1994). Reproduced using figures in Alport, Styles, and Hsieh (1994, shift-stimulus condition) 88
- Figure 25: Model data for the Alport, Styles, and Hsieh (1994) task 94
- Figure 26: Human data from van 't Wout, Lavric, and Monsell (2015). Reproduced using figures from van 't Wout, Lavric, and Monsell (2015) 98
- Figure 27: Model data for the van 't Wout, Lavric, and Monsell (2015) task. 104

ABSTRACT

Author: Fleischer, Pierson. PhD
Institution: Purdue University
Degree Received: December 2018
Title: A Unified Model of Rule-Set Learning and Selection
Committee Chair: Sébastien Hélie

The ability to focus on relevant information and ignore irrelevant information is a fundamental part of intelligent behavior. It not only allows faster acquisition of new tasks by reducing the size of the problem space but also allows for generalizations to novel stimuli. Task-switching, task-sets, and rule-set learning are all intertwined with this ability. Naturally there are many models that attempt to individually describe these cognitive abilities. However, there are few models that try to capture the breadth of these topics in a unified model and fewer still that do it while adhering to the biological constraints imposed by the findings from the field of neuroscience. Presented here is a comprehensive model of rule-set learning and selection that can capture the learning curve results, error-type data, and transfer effects found in rule-learning studies while also replicating the reaction-time data and various related effects of task-set and task-switching experiments. The model also factors in many disparate neurological findings, several of which are often disregarded by similar models